UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/749,752	12/30/2003	Matthew Mattina	42P17893	9070

59796 7590 02/02/2011

INTEL CORPORATION
c/o CPA Global
P.O. BOX 52050
MINNEAPOLIS, MN 55402

EXAMINER	
CAMPOS, YAIMA	

ART UNIT	PAPER NUMBER
2185	

NOTIFICATION DATE	DELIVERY MODE
02/02/2011	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

heather.l.adamson@intel.com

Office Action Summary	Application No.	Applicant(s)	
	10/749,752	MATTINA ET AL.	
	Examiner	**Art Unit**	
	YAIMA CAMPOS	2185	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *22 November 2010*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1 and 3-32* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1, 3-32* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)

2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)

3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.

4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.

5) ☐ Notice of Informal Patent Application

6) ☐ Other: _____.

DETAILED ACTION

1. As per the instant Application having Application number 10/749,752, the examiner acknowledges the applicant's submission of the amendment dated 7/29/2010. At this point, claims 1, 3-14, and 18 have been amended, claim 2 has been canceled and claims 21-32 have been added. Claims 1 and 3-32 are pending.

REJECTIONS BASED ON PRIOR ART

Claim Rejections - 35 USC § 103

2. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

3. Claims 1, 3-5, 7-10, 21-24, 26-29 and 32, are rejected under 35 U.S.C. 103(a) as being unpatentable over Cypher (US 2004/0010610) in view of Blake et al. (US 2004/0230751) and Jennings, III (US 6,134,631).

4. As per claim 1. **An <u>apparatus</u> for maintaining cache coherency comprising:**

a plurality of processor cores, wherein the plurality of processor cores each include a private cache; [Cypher discloses “Processing Subsystems 142” (figs. 1 and 2 and related text) where "Processing subsystems 142 may include one or more instruction and data caches" (par. 0042)]

a shared cache to be shared by the plurality of processor cores, [“Because each of processing subsystems 142 within computer system 140 may access data in memory subsystems 144,

potentially caching the data, coherency must be maintained between processing subsystems and memory subsystems 144" (par. 0042)]

wherein the shared cache includes logic, in response to receiving a write request referencing a block from a requesting processor core of the plurality of processor cores and the block not being owned, to generate a first message including an invalidation part and a write-acknowledgement part, and wherein at least the invalidate part of the first message when received by a second processor core of the plurality of processor cores is to invalidate the block in the second processor core and at least the write-acknowledgement part the first message, when received by the requesting processor core, is also to act as a write acknowledgement to the requesting processor core; and [With respect to this limitation, Cypher discloses "the invalidating request INV is a "foreign" invalidating request since it is not part of a transaction initiated by that particular device. The home memory system M also conveys the invalidating request INV to requesting device D1 (e.g., on the Multicast Network). Receipt of the INV by the requesting device indicated that shared copies have been invalidated and that write access is now allowed" (par. 0139, see pars. 0140) (thus, the invalidate message, when received by the requesting node, acting as a write acknowledgement to the requesting processor core) where "In order to gain write access, D1 initiates an RTO transaction for the coherency unit by sending an RTO request on the address network. The address network conveys the RTO request to the home memory subsystem for the coherency unit. The memory subsystem M sends an RTO response to the owning device D2. When there are non-owning active devices that have shared access to a requested coherency unit, the memory subsystem normally sends INV packets to the sharing devices (thus, the invalidate message, when received

by the non-owning sharing devices, is to invalidate the data block). However, in this example, the only non-owning sharer D1 is also the requester. Since there is no need to invalidate D1's access right, the memory subsystem may not send an INV packet to D1, thus reducing traffic on the address network. Accordingly, the memory subsystem M may return an RTO response (as opposed to a WAIT) to the requesting device D1. Upon receipt of the RTO response, D1 gains ownership of the requested coherency unit. Likewise, D2 loses ownership upon receipt of the RTO response. D1 gains write access to the requested coherency unit upon receipt of both the RTO response and the DATA packet from D2" (par. 0145) "FIG. 13A, the data packet from memory M may serve to indicate no other valid copies remain within other devices D2. In alternative embodiments, where ordering within the network is not sufficiently strong, various forms of acknowledgements (ACK) and other replies may be utilized to provide confirmation that other copies have been invalidated" (par. 0150); where the invalidate messages as taught by Cypher may be interpreted to comprise an invalidating part and a write acknowledgment part since when received by non-owning sharer nodes, the messages invalidates data and when received by requesting nodes, the messages acknowledges and allows write transactions to occur]; therefore, however Cypher does not expressly disclose **the invalidating message comprising an invalidating part and a write acknowledgement part, a ring to connect the plurality of processor cores and the shared cache, the ring to transmit the first message to the requesting processor core and second processor core, nor the particular elements taught by Cypher being in an integrated circuit**.

With respect to this limitation, **the invalidating message comprising an invalidating part and a write acknowledgement part, a ring to connect the plurality of processor cores**

and the shared cache, the ring to transmit the first message to the requesting processor core and second processor core [Blake discloses "The bus protocol set provides methods to efficiently package the various protocol constructs into a ring message so as to minimize overall coherency bus utilizations and to fit onto a small bus interface by combining the snoop command/address along with snoop responses that get ordered as the message passes through the nodes" (par. 0019) where "the nodes are interconnected by a dual concentric ring topology. Local controllers on any given node initiate bus operations on behalf of said processors and I/O adapters on that node... As the messages traverse the nodes on the ring, they trigger remote controllers to perform coherent actions such as cache accesses or directory updates. Messages arriving on each node from both directions are combined with each other and with locally generated responses to form cumulative final responses… A novel ring protocol is contemplated which efficiently packages coherency information into bus operational responses that also allow simultaneous data transfer is the direction of minimal latency" (Abstract) where read only invalidate and invalidate responses or acknowledgements are explained in (pars. 0110-0119)].

With respect to the limitation of **an integrated circuit including the elements of claim 1,** Jennings teaches a non-volatile memory with embedded programmable controller in which his plurality of modules may all implemented on a single integrated chip (storage system 50 (Fig. 1) may be a multi-chip module, or a single integrated circuit – col. 3, lines 52-58).

At the time of the invention to modify Cypher to implement the shared memory system/method taught by Cypher in a ring topology such as that taught by Blake and to further combine the coherent actions such as cache accesses or directory updates, which in the system/method taught by Cypher would includes invalidating data in foreign caches with

response which would correspond to acknowledgements taught by Cypher, as in the manner teaches combining messages/responses in a ring topology, since Blake suggests doing so would provide the benefits of efficiently packaging “coherency information into bus operational responses that also allow simultaneous data transfer is the direction of minimal latency" (Abstract) and “efficiently package the various protocol constructs into a ring message so as to minimize overall coherency bus utilizations and to fit onto a small bus interface by combining the snoop command/address along with snoop responses that get ordered as the message passes through the nodes" (par. 0019). It would have further been obvious to one of ordinary skill in the art at the time of the invention to implement the modules taught by Cypher on a single integrated circuit as taught by Jennings. By doing so, Cypher could exploit the well-known benefits of single chip integration, which includes lower manufacturing costs, and increased communication speed between the discrete elements implement on the one chip.

5. As per claim 3. **The <u>apparatus</u> of claim 1 wherein the shared cache includes one or more banks, wherein the one or more cache banks is responsible for a subset of a physical address space of the system, and wherein the block is associated with a physical address of the physical address space of the system** [Cypher discloses “Because each of processing subsystems 142 within computer system 140 may access data in memory subsystems 144, potentially caching the data, coherency must be maintained between processing subsystems and memory subsystems 144” (par. 0042) where "a domain is a group of clients that share a common physical address space” (par. 0059); thus, each of the caches comprising a portion of the total physical space of the system. Blake teaches four nodes depicted in fig. 1a, where each node comprises a System Controller Element (103), which "contains top-level cache which serves as

the central coherency point within that particular node. Both the top-level cache and the main memory are accessible by a central processor or I/O adapter within that node (104) or any of the remaining three nodes" (par. 0045); thus each cache comprising a portion the physical memory of the system].

6. As per claim 4. **The <u>apparatus</u> of claim 1 wherein the first message includes an InvalidateAndAcknowledge message, and wherein the shared cache is to generate the InvalidateAndAcknowledge message, further in response to the block being present in the shared cache and the second processor core being a custodian for the block** [Cypher teaches invalidate and acknowledge messages sent in to caches where data blocks are not owned and have the data present, which corresponds to the claimed custodian for the block since one or more caches may be in this state (pars. 145-147) and Applicant's Specification has described a custodian as merely a single processor that has a copy of the block but does not own it; see pars. 0020-0021 of Applicant's Specification; thus the embodiment in which one block is in that state corresponds to the claimed custodian. Blake further teaches "Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off" (par. 0061) thus, being present and not owned since the IM bit is off; where Blake further discloses "When MC (Multi-copy) bit is active for a particular address on a node, it indicates that one or more read-only copies of the data for this address may exist in remote caches" (par. 0056); thus, in a case where the MC bit is off, no remote caches contain copies of the data and the node caching the data in a read only state would correspond to a custodian state. Further, Blake teaches a state where MC=0, IM=1, UNOWNED BY CPS (last two states listed in fig. 2 and related text); thus, the data would be present since IM=1, unowned (or not owned) and the CP

would be the custodian since MC=0, indicating no remote processors are caching the data and the block is a custodian state; where for a “Read Only Invalidate operations, a local IM Hit always results in the cache ownership being updated to Invalid at the node… this condition must subsequently be observed at all other nodes in order to ensure that the proper cache management actions are performed at all nodes" (par. 0066)].

7. As per claim 5. **The <u>apparatus</u> of claim 1 wherein the first message includes an InvalidateAllAndAcknowledge message, and wherein the shared cache, in response to receiving the write request referencing the block from the requesting processor core of the plurality of processor cores and the block not being owned, is to generate the InvalidateAllAndAcknowledge message, further in response to the block not being present in the shared cache and none of the plurality of processor cores being a custodian for the block** [Cypher teaches invalidate and acknowledge messages sent in to caches where data blocks are not owned and have the data present, which corresponds to the claimed no custodian state for the block since one or more caches may be in this state (pars. 145-147) and Applicant's Specification has described a custodian as merely a single processor that has a copy of the block but does not own it and a no custodian state as multiple processors that have a copy of the block but do not own it; see pars. 0020-0021 of Applicant's Specification; thus the embodiment in which more than one cache is in that state corresponds to the claimed no custodian state. Blake further teaches “Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off” (par. 0061) thus, being present and not owned since the IM bit is off; where Blake further discloses “Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off” (par.

0061) where in the example, in (pars. 0105-0109; figs. 8a-8e and related text), node "N1 (801) is Read Only, IM=0, MC=1", thus indicating that the data is present, not owned since IM bit is off and that no node is a custodian node since MC bit is set, note that MC has been defined as "When MC (Multi-copy) bit is active for a particular address on a node, it indicates that one or more read-only copies of the data for this address may exist in remote caches" (par. 0056) and thus, no node would be a custodian; where for a "The Read Only Invalidate command is performed for the purpose of obtaining exclusivity of an address at the requesting node when the initial cache ownership state in the requesting node is MC=1" (par. 0099)].

8. As per claim 7. **The apparatus of claim 1 wherein the plurality of processor cores each include a merge buffer, and wherein each of the merge buffers are to coalesce multiple stores to a same block** [Cypher teaches "Each of queues… includes a plurality of entries each configured to store and address or data packet... " (pars. 0161-0162) Blake teaches "a FIFO queue for incoming messages with common addresses" (pars. 0097-0098)].

9. As per claim 8. **The apparatus of claim 1 wherein the shared cache is to fetch a second block from a memory and generate a write acknowledge message to provide a write acknowledgement to the requesting processor core in response to receiving a second write request referencing the second block, the second block not being present in the shared cache and not being owned by any of the plurality of processor cores** [Cypher teaches read to own transaction in response to a store cache miss, and appropriate responses (pars. 0052, 0134-0137; fig. 13A and related text) Blake teaches "Fetch request from a central processor which miss the top-level cache within a node will interrogate the top-level caches on the other nodes. If the fetch operation misses the top-level caches on all nodes, then the target node where the

memory address resides serves as the source for the data.... Writeback operations resulting from aged out cache data, the data is transferred directly to the target node without the need for interrogation" (par. 0046) and a "normal completion" response is generated (par. 0065) and where "an outgoing miss response... is generated for the Read Only Invalidate... as a result of the Invalid cache ownership state observed at this node" (par. 0115)].

10. As per claim 9. **The <u>apparatus</u> of claim 8 wherein the shared cache is to generate an evict message to evict a third block from an owning processor core and generate a second write acknowledge message to provide a second write acknowledgment to the requesting processor core in response to receiving a third write request referencing the third block, the third block being present in the shared cache and the owning processor core of the plurality of cores owns the third block** [Cypher teaches "An ACK data packet is a positive acknowledgement from an owning device allowing a write stream transaction to be completed" (par. 0075; see pars. 0096, 0121, 0125). Blake teaches "IM Hit… Intermediate IM Cast Out – This intermediate response is generated to signal the return of data…" (par. 0060; see par. 0066) where the cache with the IM hit would be the owner (par. 0055) and the IM hit response is sent for store type commands to grant write access to the requesting node (par. 0099; 0116)].

11. As per claim 10. **The <u>apparatus</u> of claim 1 wherein a bank of the shared cache is to be a home location for a non-overlapping portion of a physical address space associated with the block** [Cypher teaches "Each address in the address space of computer system 140 may be assigned to a particular memory subsystem 144, referred to herein as the home subsystem of the address" (par. 0043; see pars. 0052, 0053, 0055). Blake teaches node in which target memory resides (par. 0048)].

12. As per claim 21. **A method for maintain cache coherency comprising: receiving, with a shared cache, a write request referencing a block from a requesting processor core of the plurality of processor cores on a processor, wherein the plurality of processor cores each include a private cache, and wherein the plurality of cores and the shared cache are connected by a ring interconnect; generating a single message, with the shared cache, in response to receiving the write request; transmitting the single message on the ring interconnect to at least a second processor core of the plurality of processor cores and to the requesting processor core; invalidating the block in the private cache included in the second processor core in response to the second processor core receiving the single message transmitted on the ring interconnect; and write-acknowledging the write request for the requesting processor core in response to the requesting processor core receiving the single message transmitted on the ring interconnect** [The rationale in the rejection of claim 1 is herein incorporated].

13. As per claim 22. **The method of claim 21, wherein the shared cache includes one or more banks, wherein the one or more cache banks is responsible for a subset of a physical address space of a computer system including the processor, and wherein the block is associated with a physical address of the physical address space of the computer system** [The rationale in the rejection of claim 3 is herein incorporated].

14. As per claim 23. **The method of claim 21 wherein the first message includes an InvalidateAndAcknowledge message, and wherein generating the InvalidateAndAcknowledge message, with the shared cache, is further in response to the**

block being present in the shared cache and the second processor core being a custodian for the block [The rationale in the rejection of claim 4 is herein incorporated].

15. As per claim 24. **The method of claim 21 wherein the first message includes an InvalidateAllAndAcknowledge message, and wherein generating the InvalidateAllAndAcknowledge message, with the shared cache, is further in response to the block not being present in the shared cache and none of the plurality of processor cores being a custodian for the block** [The rationale in the rejection of claim 5 is herein incorporated].

16. As per claim 26. **The method of claim 21 wherein the plurality of processor cores each include a merge buffer, and wherein each of the merge buffers are to coalesce multiple stores to a same block** [The rationale in the rejection of claim 7 is herein incorporated].

17. As per claim 27. **The method of claim 21, further comprising fetching, with the shared memory, a second block from a memory and generating, with the shared memory, a write acknowledge message to provide a write acknowledgement to the requesting processor core in response to receiving a second write request referencing the second block, the second block not being present in the shared cache and not being owned by any of the plurality of processor cores** [The rationale in the rejection of claim 8 is herein incorporated].

18. As per claim 28. **The method of claim 27 further comprising generating, with the shared cache, an evict message to evict a third block from an owning processor core of the plurality of processor cores and generating a second write acknowledge message to provide a second write acknowledgment to the requesting processor core in response to receiving a third write request referencing the third block, the third block being present in**

the shared cache and the owning processor core of the plurality of cores owns the third block [The rationale in the rejection of claim 9 is herein incorporated].

19. As per claim 29. **The method of claim 21 wherein a bank of the shared cache is to be a home location for a non-overlapping portion of a physical address space associated with the block** [The rationale in the rejection of claim 10 is herein incorporated].

20. As per claim 32. **The method of claim 21 wherein the first message has a fixed deterministic latency around the ring interconnect** [Blake teaches messages transferred around the ring in a predetermined cycle or latency (par. 0018, 0053)].

21. Claims 6 and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Cypher (US 2004/0010610) in view of Blake et al. (US 2004/0230751) and Jennings, III (US 6,134,631) as applied in the rejection of claims 1 and 21 above, and further in view of Bordaz et al. (US 6,195,728) .

22. As per claim 6. The combination of Cypher, Blake and Jennings teaches **The <u>apparatus</u> system of claim 1 wherein the plurality of processor cores writes data through to the shared cache** [Cypher teaches "write-stream request initiates a transaction to allow a requesting device to write an entire coherency unit and send the coherency unit to memory… Active devices may also be configured to initiate other transaction types… using other requests" (par. 0073)]; however the combination does not expressly disclose write through .

Bordaz discloses a system/method where **the plurality of processor cores writes data through to the shared cache** as [write-through writes to reserved zones in shared cache (col. 7, lines 56-65; col. 8, line 52-col. 9, line 13)].

Cypher, Blake, Jennings and Bordaz are analogous art because they are from the same field of endeavor of computer memory and control.

At the time of the invention it would have been obvious to a person of ordinary skill in the art to modify the system and method taught by the combination of Cypher, Blake and Jennings to perform write through as taught by Bordaz. The motivation for doing so would have been because Bordaz suggests doing so would provide the benefits of [systematically updating the contents of the block since no transaction is required (col. 9, line 52-col. 9, line 13)].

Therefore, it would have been obvious to combine Cypher, Blake and Jennings with Bordaz for the benefit of creating a system/method to obtain the invention as specified in claim 6.

23. As per claim 25. **The method of claim 21 wherein the plurality of processor cores writes data through to the shared cache** [The rationale in the rejection of claim 6 is herein incorporated].

24. Claims 11 and 30 are rejected under 35 U.S.C. 103(a) as being unpatentable over Cypher (US 2004/0010610) in view of Blake et al. (US 2004/0230751) and Jennings, III (US 6,134,631) as applied in the rejection of claims 7 and 26 above, and further in view of Fletcher (US 4,445,174).

25. As per claim 11. The combination of Cypher, Blake and Jennings teaches but does not expressly disclose **The <u>apparatus</u> of claim 7 wherein each private cache of the plurality of cores are not to hold dirty data, and wherein each of the merger buffers are to hold the dirty data**.

Fletcher however teaches a multiprocessor system including a shared cache which a processor's private cache (Fig. 1, element 8) continuously stores data (permitting the merging of data (i.e. line by line) into the private memory from the main memory until an eviction is requested) –col. 1, line 62-68, and then moves the lines directly from a private cache to the shared cache, while circumventing the system's main memory (col. 2, lines 56-64).

It would have been obvious to one of ordinary skill in the art at the time of the invention for the combined teachings of Cypher, Blake and Jennings to further include Fletcher's multiprocessor system including a shared cache to his own system. By doing so, would realize improved system performance by having a means of automatically detecting lines of information moved to the shared cache, hence eliminating "ping ponging" of lines between requesting processors as taught by Fletcher in col. 2, lines 49-65.

26. As per claim 30. **The method of claim 26 wherein each private cache including in the plurality of cores are not to hold dirty data, and wherein each of the merger buffers are to hold the dirty data** [The rationale in the rejection of claim 11 is herein incorporated].

27. Claims 12-13 and 31 are rejected under 35 U.S.C. 103(a) as being unpatentable over Cypher (US 2004/0010610) in view of Blake et al. (US 2004/0230751) and Jennings, III (US 6,134,631) as applied in the rejection of claims 7 and 26 above, and further in view of Koenen (US 2004/0019891).

28. As per claim 12. The combination of Cypher, Blake and Jennings teaches **The <u>apparatus</u> of claim 1 wherein the ring is a bidirectional ring interconnect** [See fig. 1a and related text of

Blake depicting the ring as bidirectional ring interconnect]; however, the combination does not expressly disclose the bidirectional ring interconnect being **synchronous, unbuffered**.

Koenen however teaches an apparatus for optimizing performance in a multi-processing system, which includes connecting a plurality of module nodes via a synchronous, unbuffered, bi-directional ring. Referring to Fig. 1, a plurality of processing nodes (elements 12, 14 and 16) are connected for bi-directional communication (elements 12J, 14J and 16J) with the interconnect fabric (element 18). Note Koenen describes the fabric as including a ring structure in paragraph 0019, lines 9-12. The ring functions without the aid of a buffering system (i.e. unbuffered), and supports synchronous connections with a minimum static latency around the ring (paragraph 0026, lines 7-12 – the minimum latency is static).

It would have been obvious to one of ordinary skill in the art at the time of the invention, for the combined teachings of Cypher, Blake and Jennings to implement Koenen's apparatus for optimizing performance in a multi-processing system. By doing so, they would benefit by using a superior interconnection fabric (as shown by Koenen in Fig. 1, element 18) for his processing modules, which in turn could help the combination by reducing access latency and increase system performance as taught by Koenen in paragraph 0011, lines 1-15.

29. As per claim 13. **The <u>apparatus</u> of claim 12 wherein the first message has a fixed deterministic latency around the ring interconnect** [Blake teaches messages transferred around the ring in a predetermined cycle or latency (par. 0018, 0053). Konen teaches paragraph 0023 (and subsequently Table 1), describe preset latencies between each node depending on the number of nodes included in the system. With this table, the overall latency of the entire ring

interconnect is known (likewise, fixed), which allows the system to synchronize communication between nodes].

30. As per claim 31. **The method of claim 21 wherein the ring interconnect includes a synchronous, unbuffered, bidirectional, ring interconnect** [The rationale in the rejection of claim 12 is herein incorporated].

31. Claims 14 and 16-17 are rejected under 35 U.S.C. 103(a) as being unpatentable over Blake et al. (US 2004/0230751) in view of Jennings, III (US 6,134,631) and Fletcher (US 4,445,174).

32. As per claim 14. An apparatus comprising:

a plurality of cores and a shared memory connected in a ring, the shared memory to be accessible by each of the plurality of cores, wherein each of the plurality of cores includes a private memory, and a buffer [Blake teaches four nodes connected in a ring depicted in fig. 1a, where each node comprises a System Controller Element (103), which "contains top-level cache which serves as the central coherency point within that particular node. Both the top-level cache and the main memory are accessible by a central processor or I/O adapter within that node (104) or any of the remaining three nodes" (par. 0045); thus each cache comprising a portion the shared physical memory of the system; comprising “a FIFO queue for incoming messages with common addresses” (pars. 0097-0098)]

wherein the shared memory includes <u>receiving logic to receive, from a</u> <u>requesting core of the plurality of cores, a read request referencing the address, ownership logic to determine an owning processor core of the plurality of processor cores owns a block associated with</u>

the address, and [Blake discloses "Data fetch and store requests are initiated by the central processor or I/O adapters, and are processed by the local controllers contained within the SCE (103)" (par. 0045, see par. 0046) where "As the operation passes through each remote node, remote fetch controllers interrogate the top-level cache on that remote node and perform any necessary system coherency actions" (par. 0048) where "For Read Only, Fetch Exclusive... a local IM Hit condition always results in cache data being sourced from the node" (par. 0066)] **eviction logic coupled to the receiving logic and the ownership logic, the eviction logic to generate an evict message referencing the address and the owning processor core in response to the receiving logic receiving the read request and the ownership logic determining the owning processor core owns the block** [Blake discloses "IM Hit... Intermediate IM Cast Out-This intermediate response is generated to signal the return of data when the IM bit is on..." (par. 0060) where "For Read Only, Fetch Exclusive... a local IM Hit condition always results in cache data being sourced from the node" (par. 0066)].

Blake does not expressly disclose **an integrated circuit including:** the modules of claim 14, nor **a merge buffer to purge data to the shared memory**.

With respect to the limitation of **an integrated circuit including the elements of claim 14,** Jennings teaches a non-volatile memory with embedded programmable controller in which his plurality of modules may all implemented on a single integrated chip (storage system 50 (Fig. 1) may be a multi-chip module, or a single integrated circuit – col. 3, lines 52-58).

Fletcher teaches a multiprocessor system including a shared cache which a processor's private cache (Fig. 1, element 8) continuously stores data (permitting the merging of data (i.e. line by line) into the private memory from the main memory until an eviction is requested) –col.

1, line 62-68, and then moves the lines directly from a private cache to the shared cache, while circumventing the system's main memory (col. 2, lines 56-64).

It would have further been obvious to one of ordinary skill in the art at the time of the invention to implement the modules taught by Blake on a single integrated circuit as taught by Jennings. By doing so, Blake could exploit the well-known benefits of single chip integration, which includes lower manufacturing costs, and increased communication speed between the discrete elements implement on the one chip; and to further modify the combined teachings of Blake and Jennings include Fletcher's multiprocessor system including a shared cache to his own system. By doing so, would realize improved system performance by having a means of automatically detecting lines of information moved to the shared cache, hence eliminating "pingponging" of lines between requesting processors as taught by Fletcher in col. 2, lines 49-65.

33. As per claim 16. **The apparatus of claim 14, wherein the shared memory is a shared cache including a plurality of blocks, and wherein the shared cache is capable of holding each of the plurality of blocks in a cache coherency state** [Blake teaches shared memory (fig. 1a and related text; pars. 0043, 0046) plurality of coherency states (pars. 0057-0065; 0105-0109)].

34. As per claim 17. **The apparatus of claim 16, wherein the cache coherency state for each of the plurality of blocks is selected from a group consisting of (l) a not present state, (2) a present and owned by a core of the plurality of cores state, (3) a present, not owned, and custodian is a core of the plurality of cores state, and (4) a present, not owned, and no custodian state** [Blake discloses "Miss-This local response is generated if the cache ownership state at the node is found to be invalid" (par. 0059); "IM Hit-This local response is generated at a

node if the IM bit is on..." (par. 0060) where "When the IM (Intervention Master) bit is active for a particular address on a node, it indicates that this node was the most recent to cache in new data and receive cache ownership for that address" (par. 0055); "Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off" (par. 0061) thus, being present and not owned since the IM bit is off; where Blake further discloses "When MC (Multi-copy) bit is active for a particular address on a node, it indicates that one or more read-only copies of the data for this address may exist in remote caches" (par. 0056); thus, in a case where the MC bit is off, no remote caches contain copies of the data and the node caching the data in a read only state would correspond to a custodian state. Further, Blake teaches a state where MC=0, IM=1, UNOWNED BY CPS (last two states listed in fig. 2 and related text); thus, the data would be present since IM=1, unowned (or not owned) and the CP would be the custodian since MC=0, indicating no remote processors are caching the data; "Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off" (par. 0061) where in the example, in (pars. 0105-0109; figs. 8a-8e and related text), node "N1 (801) is Read Only, IM=0, MC=1", thus indicating that the data is present, not owned since IM bit is off and that no node is a custodian node since MC bit is set, note that MC has been defined as "When MC (Multi-copy) bit is active for a particular address on a node, it indicates that one or more read-only copies of the data for this address may exist in remote caches" (par. 0056)].

35. Claim 15 is rejected under 35 U.S.C. 103(a) as being unpatentable over Blake et al. (US 2004/0230751) in view of Jennings, III (US 6,134,631) and Fletcher (US 4,445,174) as applied in the rejection of claim 14 above and further in view of Koenen (US 2004/0019891).

36. As per claim 15. **The apparatus of claim 14, wherein the ring includes a bi-directional ring interconnect** [Blake teaches bidirectional ring interconnect (fig. 1a and related text)]; however, the combination does not expressly disclose the ring as **synchronous unbuffered**.

Koenen however teaches an apparatus for optimizing performance in a multi-processing system, which includes connecting a plurality of module nodes via a synchronous, unbuffered, bi-directional ring. Referring to Fig. 1, a plurality of processing nodes (elements 12, 14 and 16) are connected for bi-directional communication (elements 12J, 14J and 16J) with the interconnect fabric (element 18). Note Koenen describes the fabric as including a ring structure in paragraph 0019, lines 9-12. The ring functions without the aid of a buffering system (i.e. unbuffered), and supports synchronous connections with a minimum static latency around the ring (paragraph 0026, lines 7-12 – the minimum latency is static).

It would have been obvious to one of ordinary skill in the art at the time of the invention, for the combined teachings of Blake, Jennings and Fletcher to implement Koenen's apparatus for optimizing performance in a multi-processing system. By doing so, they would benefit by using a superior interconnection fabric (as shown by Koenen in Fig. 1, element 18) for his processing modules, which in turn could help the combination by reducing access latency and increase system performance as taught by Koenen in paragraph 0011, lines 1-15.

37. Claims 18-19 are rejected under 35 U.S.C. 103(a) as being unpatentable over Blake et al. (US 2004/0230751) in view of Bordaz et al. (US 6,195,728) and Jennings, III (US 6,134,631).

38. As per claim 18. Blake A system comprising:

a plurality of cores and a shared memory to be coupled together with a bi-directional ring interconnect, the shared memory is to be accessible by each of the plurality of cores, and the shared memory is to include a plurality of blocks, [Blake teaches a shared memory symmetrical multiprocessing system (par. 0043) comprising four nodes connected in a bi-directional ring depicted in fig. 1a, where each node comprises a System Controller Element (103), which "contains top-level cache which serves as the central coherency point within that particular node. Both the top-level cache and the main memory are accessible by a central processor or I/O adapter within that node (104) or any of the remaining three nodes" (par. 0045)]

each of the plurality of blocks capable of being held by logic in the shared memory in a not present state; [Blake discloses "Miss-This local response is generated if the cache ownership state at the node is found to be invalid" (par. 0059)]

a present and owned by a core of the plurality of cores state; [Blake discloses "IM Hit-This local response is generated at a node if the IM bit is on..." (par. 0060) where "When the IM (Intervention Master) bit is active for a particular address on a node, it indicates that this node was the most recent to cache in new data and receive cache ownership for that address" (par. 0055)]

a present, not owned, and a core of the plurality of cores is a custodian state; and [Blake discloses "Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off" (par. 0061) thus, being present and not owned since the

IM bit is off; where Blake further discloses "When MC (Multi-copy) bit is active for a particular address on a node, it indicates that one or more read-only copies of the data for this address may exist in remote caches" (par. 0056); thus, in a case where the MC bit is off, no remote caches contain copies of the data and the node caching the data in a read only state would correspond to a custodian state. Further, Blake teaches a state where MC=0, IM=1, UNOWNED BY CPS (last two states listed in fig. 2 and related text); thus, the data would be present since IM=1, unowned (or not owned) and the CP would be the custodian since MC=0, indicating no remote processors are caching the data]

a present, not owned, and no core of the plurality of cores is a custodian state; and [Blake discloses "Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off" (par. 0061) where in the example, in (pars. 0105-0109; figs. 8a-8e and related text), node "N1 (801) is Read Only, IM=0, MC=1", thus indicating that the data is present, not owned since IM bit is off and that no node is a custodian node since MC bit is set, note that MC has been defined as "When MC (Multi-copy) bit is active for a particular address on a node, it indicates that one or more read-only copies of the data for this address may exist in remote caches" (par. 0056)]

a system memory associated with the processor to hold elements to be stored by the shared memory [Blake teaches "The system's main memory is distributed across the nodes… hardware mapping registers… determines if the main memory location exists on that node" (par. 0049)].

Blake does not expressly disclose **a processor or single circuit including**: the elements of claim 18, **wherein each of the plurality of cores is to be associated with a private cache memory,** nor the ring as an **unbuffered** ring.

Bordaz discloses a shared memory multiprocessing system where each of the plurality of cores is associated with a private cache memory and connecting in an unbuffered ring configuration as [caches within processors and ring configuration which is shown as simple links without including a buffer (fig. 1 and related text)].

Jennings teaches a non-volatile memory with embedded programmable controller in which his plurality of modules may all implemented on a single integrated chip (storage system 50 (Fig. 1) may be a multi-chip module, or a single integrated circuit – col. 3, lines 52-58).

At the time of the invention it would have been obvious to one having ordinary skill in the art to modify Blake to include private caches in each of the processing cores and have the ring as an unbuffered ring, as taught by Bordaz, since doing so would provide the Benefits of fast access speed to data in the processors' private caches as well as fast transfer speed on the ring. It would have further been obvious to one of ordinary skill in the art at the time of the invention to implement the modules taught by the combination of Blake and Bordaz on a single integrated circuit as taught by Jennings, since doing so would provide the well known benefits of single chip integration, which includes lower manufacturing costs, and increased communication speed between the discrete elements implement on the one chip.

39. As per claim 19. **The system of claim 18, wherein each of the plurality of blocks is a home location for a subset of a physical address space [Blake teaches node in which target memory resides** (par. 0048-0048)**]**.

40. Claim 20 is rejected under 35 U.S.C. 103(a) as being unpatentable over Blake et al. (US 2004/0230751) in view of Bordaz et al. (US 6,195,728) and Jennings, III (US 6,134,631) and further in view of Cypher (US 2004/0010610).

41. As per claim 20. **The system of claim 19, wherein the shared cache is to generate a first message to invalidate a requested block in all cores of the plurality of cores except for a requesting core of the plurality of cores, in response to receiving a write request referencing the requested block from the requesting core and requested block being held in the present, not owned, and no core of the plurality of cores is a custodian state** [Blake further teaches "Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off" (par. 0061) thus, being present and not owned since the IM bit is off; where Blake further discloses "Read Only Hit-This local response is generated at a node if the cache ownership state is found Read Only and the IM bit is off" (par. 0061) where in the example, in (pars. 0105-0109; figs. 8a-8e and related text), node "N1 (801) is Read Only, IM=0, MC=1", thus indicating that the data is present, not owned since IM bit is off and that no node is a custodian node since MC bit is set, note that MC has been defined as "When MC (Multi-copy) bit is active for a particular address on a node, it indicates that one or more read-only copies of the data for this address may exist in remote caches" (par. 0056) and thus, no node would be a custodian; where for a "The Read Only Invalidate command is performed for the purpose of obtaining exclusivity of an address at the requesting node when the initial cache ownership state in the requesting node is MC=1" (par. 0099) where invalidate at the requesting node is not performed since the requesting node will obtain exclusive ownership of

the data] but does not expressly disclose generating a first message to invalidate the requested block in all cores of the plurality of cores.

With respect to the limitation generating a first message to invalidate the requested block in all cores of the plurality of cores, except to the requesting core Cypher discloses [invalidate messages sent in to caches where data blocks are not owned and have the data present, which corresponds to the claimed no custodian sate for the block since one or more caches may be in this state (pars. 145-147, see pars. 0139-0141) and Applicant's Specification has described a custodian as merely a single processor that has a copy of the block but does not own it and a no custodian state as multiple processors that have a copy of the block but do not own it; see pars. 0020-0021 of Applicant's Specification; thus the embodiment in which more than one cache is in that state corresponds to the claimed no custodian state.

At the time of the invention, it would have been obvious to one having ordinary skill in the art to modify the combination of Blake, Bordaz and Jennings to generate a first message to invalidate the requested block in all cores of the plurality of cores except the requesting cores being in a no custodian state as taught by Cypher, since doing so would provide the benefits of [allowing requesting device exclusive access rights while transitioning responsibilities and access rights correctly (par. 0142) and reducing traffic on the network (par. 0145)].

ACKNOWLEDGMENT OF ISSUES RAISED BY THE APPLICANT

Response to Amendment

42. Applicant's arguments filed on 11/22/2010 have been fully considered but are moot in view of the new ground(s) of rejection.

43. However, some of Applicant's arguments were not deemed persuasive.

44. As required by **M.P.E.P. § 707.07(f)**, a response to these arguments appears below.

ARGUMENTS CONCERNING PRIOR ART REJECTIONS

45. Applicant arguments regarding intended use recitations, for example for of the claim language including "logic to", "to generate a first message", "capable of being held by logic in the shared memory", stating that the claim language eliminates intended use are not deemed persuasive since the recited language does not change the scope of the claim enough to eliminate intended use. A recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use then it meets the claim. Further, a certain structure being "capable" of performing a certain functionality does not require that the structure perform the listed functionality but merely that it not be expressly precluded from doing so.

46. For example, amending the system claims to recite a system or structure <u>configured to</u> perform a certain functionality such as the recited structures <u>configured to… include logic</u>, <u>configured to generate</u>, <u>configured to</u> hold by logic in the shared memory would effectively eliminated intended use language from the claims since the structures would be positively recited as being <u>configured to</u> perform the listed functionality.

CLOSING COMMENTS

a. STATUS OF CLAIMS IN THE APPLICATION

47. The following is a summary of the treatment and status of all claims in the application as recommended by **M.P.E.P. 707.07(i)**:

a(1) CLAIMS NO LONGER UNDER CONSIDERATION

48. Claim 2 has been canceled.

a(2) CLAIMS REJECTED IN THE APPLICATION

49. Per the instant office action, claims 1 and 3-32 have received an action on the merits and are subject of a non-final rejection.

b. DIRECTION OF FUTURE CORRESPONDENCES

50. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Yaima Campos whose telephone number is (571) 272-1232, and email address is yaima.campos@uspto.gov . The examiner can normally be reached on Monday to Friday 8:30 AM to 5:00 PM.

51. If attempts to reach the above noted Examiner by telephone are unsuccessful, the Examiner's supervisor, Mr. Sanjiv Shah, can be reached at the following telephone number: Area Code (571) 272-4098.

The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions

on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

January 28, 2011

/Yaima Campos/
Examiner, Art Unit 2185